solarwinds

HOW CONFIGURATION MANAGEMENT CAN SUPERCHARGE YOUR NETWORK PERFORMANCE MONITORING

By: Craig Tobias, CEO of Tobias International

SOLARWINDS
WHITEPAPER

HOW CONFIGURATION MANAGEMENT CAN SUPERCHARGE YOUR NETWORK PERFORMANCE MONITORING

EXECUTIVE SUMMARY

The purpose of any network monitoring system is to ensure the highest level of network availability and performance. This document outlines a best practices approach to using SolarWinds® Network Performance Monitor (NPM) and Network Configuration Manager (NCM) as a crucial part of a comprehensive strategy for effective network monitoring and management.

SITUATION

To understand this objective, we must first understand three key functional areas of network management. We also need to know how SolarWinds NPM and NCM address these areas. Network management comprises the following:

- Fault management.
- Performance management.
- Configuration management.

The relationship between network management and SolarWinds NCM and NPM looks like this:

SolarWinds Orion® Module	Functional Area	Items Managed
NPM	Fault Management	• Detection of link or circuit outages and errors. • Detection of equipment failures. • Detection of major data disruptions.
NPM, NTA	Performance Management	• Collection of network statistics. • Display of equipment health. • Display of link or circuit health.
NCM	Configuration Management	• Configuration reporting. • Backup of device configurations. • Validation for configuration standards.

You can see that the success of a fault management system is dependent on the proper configuration of network device management features, network management software, and the host systems they reside on. All must be scaled according to the size of the network and its projected growth.

Configuration management is the functional area that addresses the backup, maintenance, standards, and compliance of network equipment configurations. Configuration management is just as critical as fault and performance management in terms of ensuring a reliable, highly available network.

THE PROBLEM

Studies have shown that nearly 80% of all network outages in an organization are related to human error, while only a small number of outages are related to equipment failure, such as a faulty power supply or cable. *Gartner: Ronni J. Colville and George Spafford. Configuration Management for Virtual and Cloud Infrastructures.*

However, configuration management is usually only addressed by most organizations after fault and performance management solutions have been deployed. There are practical as well as psychological reasons for this. While configuration management will eventually have the greatest impact on overall network availability, it requires the most planning and is the least tangible. Deploying a fault and performance management system, such as NPM, gives IT pros immediate visibility into the network. Doing so also measures the positive impact that configuration management has once it is deployed. However, to receive the full benefits of comprehensive network management, configuration management should be viewed as being equal with performance and fault management.

THE SOLUTION

Deploying SolarWinds NCM in your environment helps remove the possibility of human error through a number of mechanisms, including:

- Initial provisioning of devices to standards.
- Carefully managing configuration changes.
- Monitoring configurations for unwanted changes.
- Auditing configurations for compliance.
- Troubleshooting network performance and faults.
- Archiving configurations for disaster recovery.

SolarWinds NCM can reduce the potential for misconfiguration by using automated device configuration tools to build the network. For example, SolarWinds NCM can be used to build the initial configuration before integrating it into the provisioning process. This expedites the deployment of configurations to devices, and helps ensure that the configuration adheres to the proper standards established by the organization.

This is accomplished by first setting up the device with an IP address, SNMP, and login credentials. Next, declare the device to NCM so that it can push the complete operational configuration of the device. This can be done in a staging area, or when the device is able to communicate with SolarWinds.

Once a configuration has been deployed and meets the organization's standards, it then becomes critical to ensure that those configurations remain standardized.

For example, let's say an organization has a large enterprise network that uses the Cisco® 3750 series. In addition, they have standardized the 3750 configuration so that the PortFast BPDU Guard is enabled on all access layer ports. This prevents someone connecting to a switch or hub at their desk from affecting the topology of the spanning tree. However, during the course of troubleshooting, an engineer turns off this feature on all ports to determine if the problem occurs when attempting to connect to a specific brand of interface card. Even though this wasn't

the underlying issue, the PortFast BPDU Guard was inadvertently left off. While this may not cause an immediate setback, it would be problematic if someone connected a consumer-grade network switch or hub at their desk.

To prevent these types of problems, NCM includes policy auditing and remediation, where standardization rules can be established to report on and reinforce configuration policies. In this case, it would automatically re-enable the Portfast BPDU guard to all access ports when a policy audit discovered the feature had been disabled.

Finally, SolarWinds NCM provides the capability to alert and report non-compliant changes. Almost all enterprise-grade network equipment will send either a syslog or SNMP trap when a change is made to the configuration file. SolarWinds can receive these messages and generate a time-based alert if the change notification was received outside normal change windows. Simultaneously, NCM can collect the new configuration file and determine exactly what was modified. In doing so, management in charge of change control can review the modification, determine who made the modification, understand why the modification was made, and know exactly what was modified.

For a more detailed look at how NCM and NPM work together to deliver comprehensive performance, fault, and configuration management, consider the following examples and case studies.

NETWORK TUNNELS

Over the years, networking has evolved as the complex demands on networks have increased. For technical as well organizational reasons, a department may require tunnels to be configured to route specific types of traffic to different locations. These tunnels are often configured without much planning to meet a specific need. Often, engineers forget to specify the tunnel speed because they are configuring this on a device with 1GB or 10GB interfaces, and then wrongly assume the tunnel speed is dynamic, when it is not. If a speed is not specified, the default speed is used. For example, a GRE tunnel on a Cisco device will default to 8mb. For most situations, this is plenty. However, if this tunnel is used for video streaming or nightly backups, it might not be enough. This is where SolarWinds NPM and NCM working together can assist to identify and resolve the problem.

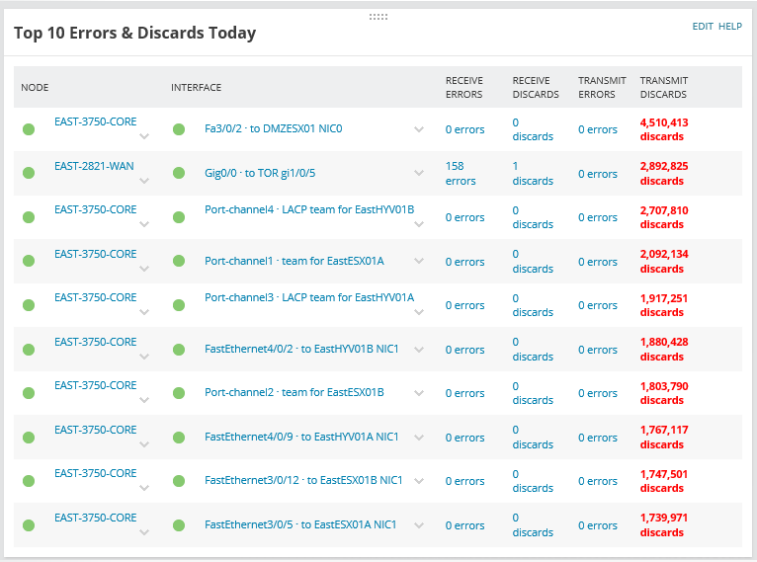

Top 10 Errors & Discards Today

EDIT HELP

NODE	INTERFACE	RECEIVE ERRORS	RECEIVE DISCARDS	TRANSMIT ERRORS	TRANSMIT DISCARDS
EAST-3750-CORE	Fa3/0/2 · to DMZESX01 NIC0	0 errors	0 discards	0 errors	4,510,413 discards
EAST-2821-WAN	Gig0/0 · to TOR gi1/0/5	158 errors	1 discards	0 errors	2,892,825 discards
EAST-3750-CORE	Port-channel4 · LACP team for EastHYV01B	0 errors	0 discards	0 errors	2,707,810 discards
EAST-3750-CORE	Port-channel1 · team for EastESX01A	0 errors	0 discards	0 errors	2,092,134 discards
EAST-3750-CORE	Port-channel3 · LACP team for EastHYV01A	0 errors	0 discards	0 errors	1,917,251 discards
EAST-3750-CORE	FastEthernet4/0/2 · to EastHYV01B NIC1	0 errors	0 discards	0 errors	1,880,428 discards
EAST-3750-CORE	Port-channel2 · team for EastESX01B	0 errors	0 discards	0 errors	1,803,790 discards
EAST-3750-CORE	FastEthernet4/0/9 · to EastHYV01A NIC1	0 errors	0 discards	0 errors	1,767,117 discards
EAST-3750-CORE	FastEthernet3/0/12 · to EastESX01B NIC1	0 errors	0 discards	0 errors	1,747,501 discards
EAST-3750-CORE	FastEthernet3/0/5 · to EastESX01A NIC1	0 errors	0 discards	0 errors	1,739,971 discards

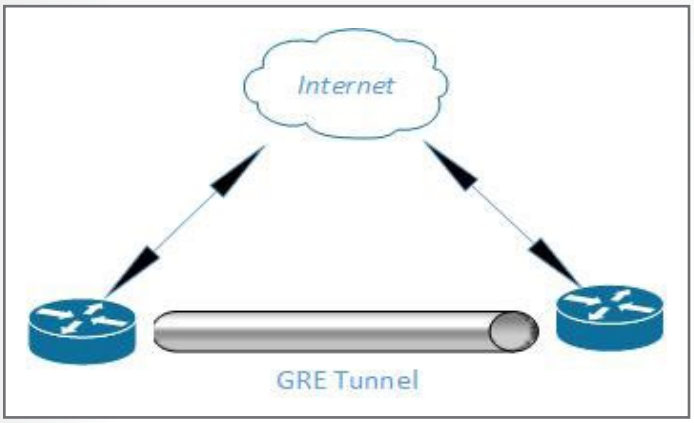

Once NPM and NC are installed, the network summary page on NPM lists interfaces with the highest utilization throughput. This allows many organizations to see the time issue for the first time. Based on this, the engineer can pull up the configuration in NCM and instantly see that no speed was set for the tunnels when they were created. Once these high-utilization tunnels are identified, the engineer can understand why certain end-users complain about system slowness or poor VoIP call quality due to jitter and other IPSLA-related problems.

Armed with this knowledge, the engineer can then use NCM to help ensure that speeds are set on all tunnels moving forward. Furthermore, they can use NCM to prevent the issue from happening again, because NCM has the ability to detect and even remediate out-of-process changes made to GRE tunnels and the network.

INTERFACE ERRORS

Network engineers have the ability to access the same data and counters from the command line as SolarWinds does via SNMP. However, when an engineer accesses this data from the command line, they are seeing a snapshot of this data out of context. From the command line, an engineer might see hundreds of errors on an interface, but wouldn't be able to tell when those errors occurred. It is impractical, if not impossible, to find errors by accessing an interface, especially on a network of significant size.

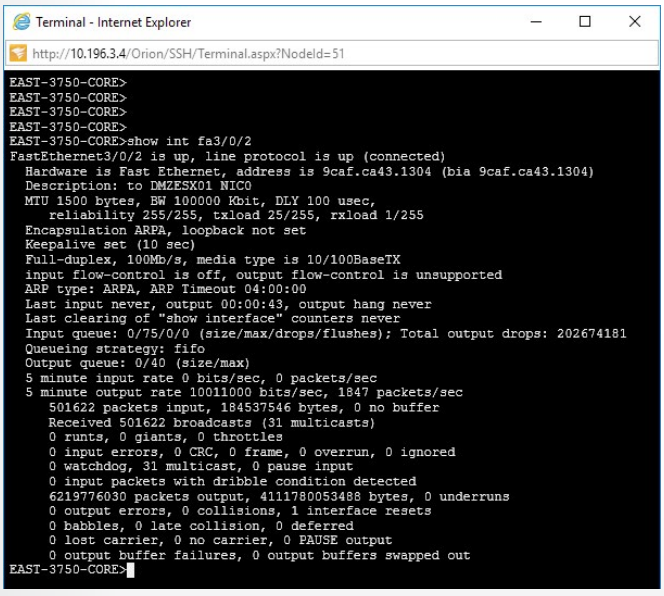

NPM periodically can access an interface (the default is every 15 minutes), and collect throughput, errors, and other counters across the entire network. In doing so, it provides a simple-to-read display that allows network engineers to quickly determine if an interface on a network device is having problems. NPM has taken this ability a step further, too, by providing a dashboard that reveals the top 10 interface errors across the network. This gives engineers the ability to quickly identify the interface that is most likely contributing to the degradation of overall network performance.

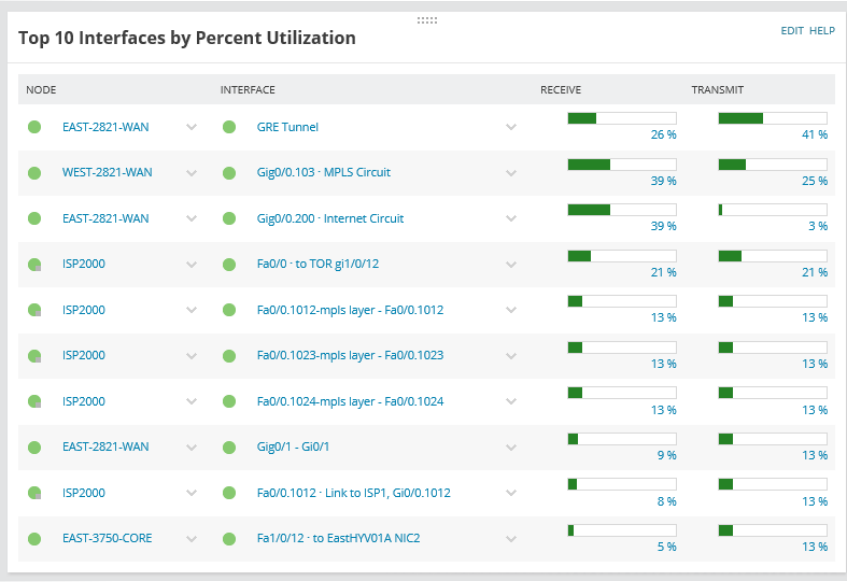

Network interfaces can have errors for a number of reasons depending on the interface type. These reasons may include a bad cable or a failing interface. However, the most common cause of network interface errors is a configuration mismatch either on the local interface or on the remote device. These misconfigurations can include speed mismatch, duplex mismatch, or auto-negotiate on one end, but not the other. Once an interface has been identified as having significant ongoing errors, the engineer can use NCM to access the configuration of the identified interface and the remote interface to determine if there are any configuration mismatches. If an error is identified, the error can be fixed in NCM, and the new configuration can be pushed out to the misconfigured device, which will resolve the problem.

LOST BGP NEIGHBOR

Border Gateway Protocol (BGP) is a WAN routing protocol often used by organizational sites or campuses. BGP should not re-converge unless there has been a network change on a network that has been operating for any length of time. BGP might re-converge if there has been a loss of a routing neighbor, which could signify the start of a major network outage. Network devices can be configured to send an SNMP trap to NPM if a BGP neighbor is lost.

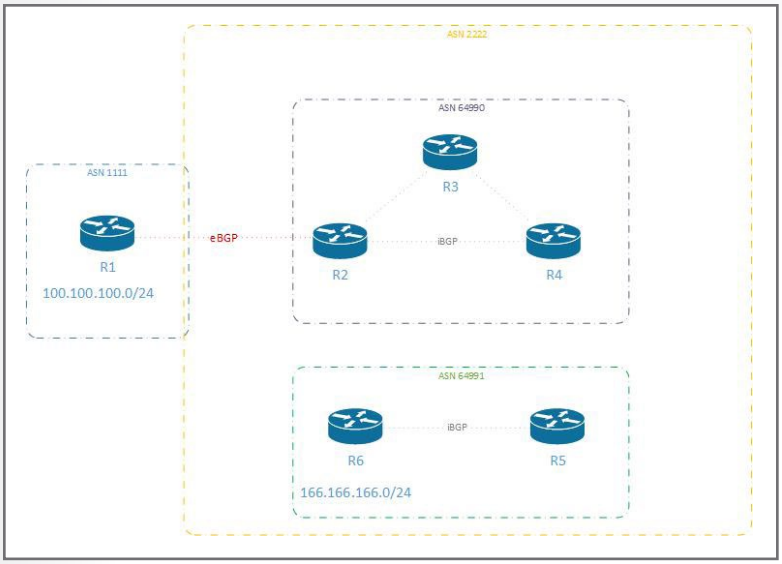

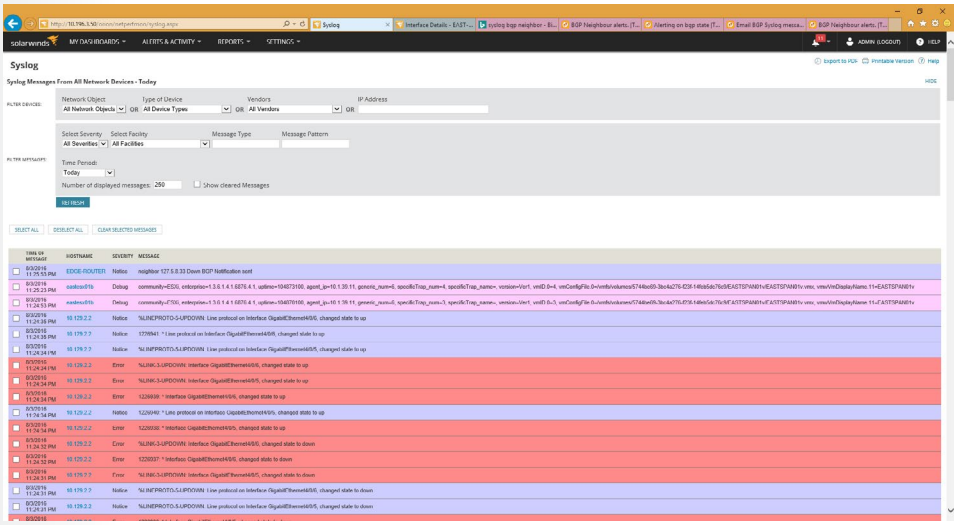

NPM should be configured to alert if a lost BGP neighbor trap is received. Once an alert is received, the network engineer should immediately refer to NCM to see if another engineer could have made a configuration change prior to the loss of that BGP neighbor.

If the loss was due to a configuration change, the last good configuration stored in NCM should be reloaded in the router to restore service.

CASE STUDY: WEB CONTENT DELIVERY

Now that we've gone over some examples, let's look at a few case studies.

Let's say you own a baked goods business called Yuma Nomnom. Yuma Nomnom doesn't actually bake pastries. It has a contract with a large convenience store chain and works with 140 nationwide bakeries that deliver and stock those chain stores with baked goods. You simply transfer the orders from the convenience stores to the appropriate bakeries. Each order must

be received by 2:00pm every day, and delivered to the various bakeries by 6:00pm every evening. This schedule ensures that the orders are processed on time.

The order from the convenience store starts to file transfer each day at 1:00pm, and normally takes about 30 minutes to complete. However, for the past three days, the order has been taking 56 minutes to complete the transfer.

The first thing you check is file size to see if it has increased, but it hasn't. The next thing you do is check to see if you are having issues receiving files from any other clients, and you aren't. Now you check NCM to see if there have been any config changes made to the gateway or firewall in the past week, and there haven't been. Next, you call the IT director at the data center used by the convenience store chain to send files to see if they've made any changes, but you are told they haven't.

Finally, you decide to use NPM NetPath™ to set up a path between Yuma Nomnom and the client's data center. You notice that everything appears to be working well, until you see that a service provider called Arctic Gecko has three routers in the path, and four paths through their network. You also see that the link between the second and third routers in the path are experiencing unusually high latency times.

Armed with this information, you open a ticket with Arctic Gecko to resolve the issue. Later, you find out that Arctic Gecko had make changes to their network and had a routing loop that was causing packets to bounce between their routers, consuming large amounts of bandwidth, slowing traffic through these devices.

CASE STUDY: 4G BACKUP LINK

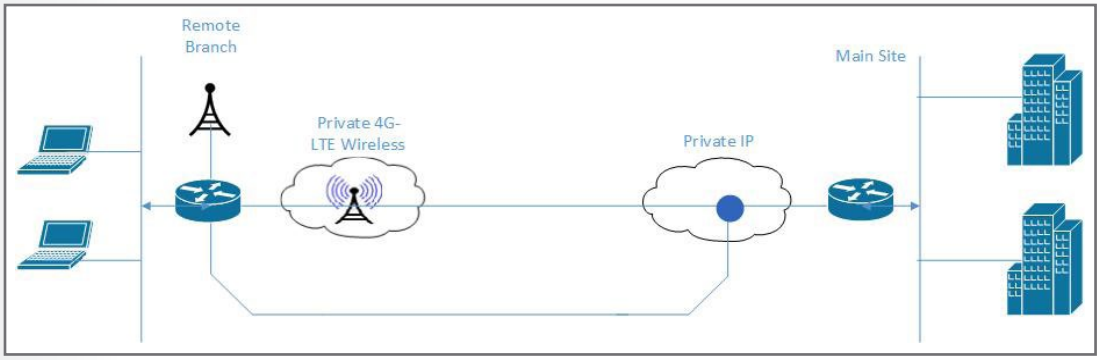

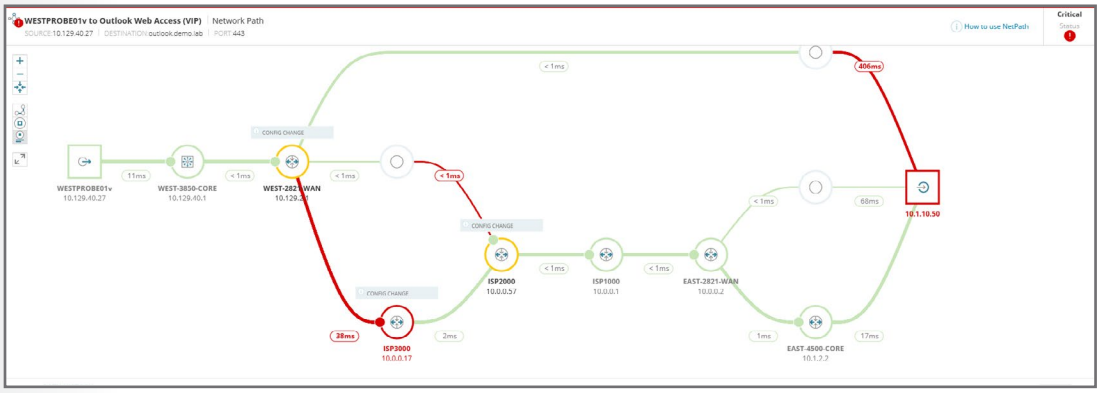

Now let's look at another case study involving Yuma Nomnom.

It is a few years later, and you've expanded the business to include five retail sites. A router at each site communicates with the corporate office using an IPSEC tunnel. The router at each retail site is configured with a primary IP connection through the local cable ISP, as well as a 4G wireless backup interface. There is also a server at each site that is used for point of sale, credit card authorization, ordering products, maintaining inventory, and tracking employee hours for payroll. If data communication to the corporate office is lost, the store cannot sell any products. To ensure optimal data connectivity, each site router is connected to the local cable ISP as the primary means of communication. Should the primary connection fail, the 4G wireless will sync up to the wireless network to maintain IP connectivity with the corporate office. Sending data over the 4G network is far more expensive because the wireless provider charges by the megabyte, but at least the store can failover and continue to operate until the primary cable ISP service can be restored. Once service is restored, the router will automatically disconnect from the 4G network and use the local cable ISP.

The general manager of several of your stores has noticed that it is taking an unusually long time for credit cards to process, products to be ordered, and confirmations for successfully placed orders to come through.

To better understand the issue, you configure all five servers to act as a NetPath probe for visibility from the retail site back to the corporate offices. You quickly see that the two sites reporting slowdowns are sending all of their data over the 4G network, not the local cable ISP, even though the local cable ISP is fully operational.

After investigating further, you see that a few weeks prior, each site lost connectivity to the local cable ISP and switched over to 4G, but never switched back once cable ISP service was restored. This was due to a misconfiguration on the router, and there was no NCM policy audit to detect or automatically remediate the config error. Fortunately, NetPath was able to catch the issue, allowing performance to be restored.

To prevent future issues, you set up an alert in NPM to manually check the circuit in case the wireless interface is up for more than an hour. In addition, you create an escalation for verifications every two hours after the initial check. As a final step, you use an NCM policy template to make sure all the stores were configured correctly, allowing the router to disconnect when local cable ISP service is restored.

SUMMARY

When planning and implementing a network management solution, it is important to understand fault, performance, and configuration management. But to achieve the greatest benefit, you must also understand how these functions work together to enhance the effectiveness of the entire network. In other words, network management is a perfect example of a whole being greater than the sum of its parts.

ABOUT THE AUTHOR

Craig Tobias, CEO of Tobias International, has been designing, building, and deploying management solutions for 25 years. He is passionate about helping individuals and organizations navigate the complexities of setting up robust and effective enterprise monitoring and operations. Tobias International specializes in deploying and integrating SolarWinds products.